RIT Business Analytics Competition Spring 2023

University of Rochester, Simon Business School

Yitong Li / Business Analytics / yli295@simon.rochester.edu

Yuxiao Liao / Business Analytics / yliao21@simon.rochester.edu

Zhiwei Liang / Marketing Analytics / zliang19@simon.rochester.edu

Tingting Liu / Marketing Analytics / tliu44@simon.rochester.edu

Xinyu Su / Marketing Analytics / xsu9_simon@simon.rochester.edu

Executive Summary

After thoroughly analyzing RIT-Commerce's identified fake reviews, we have designed a predictive model for the company to sort out potential fake reviews to address the bias caused by those untruthful comments. This report outlines the various methodologies, tools, and techniques that we utilized to achieve such goals and to provide RIT-Commerce with actionable insights that can help enhance its business operations. Whether it's refining its lending criteria or implementing more robust risk management strategies, there are several key areas where RIT-Commerce can make tangible improvements that will be helpful to improve the authenticity of its review section.

RIT-Commerce focuses on designing a fake review detection model to reduce bias within its consumer community. Since fake reviews are generated by actual users, our model takes into account many aspects of every single comment by applying the sentiment analysis method. This is because positive fake reviews would usually differ from real reviews on sentiment intensity and subjectivity.

The model was created using ANN, an advanced machine-learning ensemble model. We integrated multiple methods because we made the possible effort to balance the accuracy and flexibility trade-off while also considering the overfitting problem (Brownlee). After a complex tuning process, we completed this model with an F1 score of 0.39 yielded. Applying this model would help RIT-Commerce to better filter out biased reviews.

Since RIT-Commerce would like to discover the relationship between fake reviews and overall platform performance and profitability, we used a regression model to find out this correlation. It turns out that the amount of fake reviews has a notable impact on the sales and helpfulness of the review section, in addition to the reputation and accountability of the overall brand image.

Data Preparation

The train_review dataset contains 392,425 rows of review records and 11 variables, with "fake_review" serving as the target variable for default prediction. The test_review dataset contains 43,461 rows of review records with the same variables as the training set, except for "fake_review".

To optimize the training dataset for better prediction, several actions were taken. Firstly, rows with misaligned data were removed, resulting in a loss of 550 rows that had a negligible impact on the training dataset. All numeric variables were converted into "int" format to ensure further predictive analysis, and the "review_date" variable was divided into three sub-variables representing Year, Month, and Day. Additionally, "is_weekday" was created to indicate whether the review occurred on a weekday or not. "Label_productid" was created by multiplying "fake_asin" and "product_id" to label the product that had fake reviews before. Two text variables, "review_title" and "review_body," were merged into a single "review" feature to enable complete text sentiment analysis.

Before conducting the sentiment analysis, text preprocessing was conducted in three steps, including symbol removal (gibberish, punctuation), stop-word removal, and word stemming. Following the text cleaning, two common emotional indices of text, polarity and subjectivity, were generated as new data features. Polarity reflects the degree of positivity or negativity expressed in the text and is often measured on a scale of -1 to +1, where -1 represents a very negative sentiment, +1 represents a very positive sentiment, and 0 represents a neutral sentiment. Subjectivity reflects the degree to which the text expresses a personal opinion, feeling, or emotion, and is typically measured on a scale of 0 to 1, where 0 represents a completely objective statement and 1 represents a completely subjective statement.

Features with IDs are removed as they do not have exact meanings to our model. After the data preparation, the training set comprised {review_rating, number_of_photos, helpful_vote, review, year, month, day, is_weekday, polarity, subjectivity, labeled_product_id (get dummies)} as the existing analysis features.

Data Analysis

At the beginning of our data analysis for the dataset, we utilized Python to conduct the analysis directly. We first examined the fake reviews and discovered that the dataset contained a total of 392,760 reviews, of which 372,695 were genuine and 20,065 were fake. The ratio between genuine and fake reviews was 19:1, leading us to conclude that this dataset is highly imbalanced (Figure 1). As a result, when handling model parameters, we need to adjust our approach to address this imbalance.

Next, we conducted a feature analysis using a logistic regression model (Figure 2). We found that review rating, number of photos, year, month, is_weekday, polarity, and subjectivity had a significant correlation with fake reviews. This conclusion is supported by the Logit regression results shown in Figure 2, where the P-values are all less than 0.05, and the coefficients exhibit high values. In particular, polarity and subjectivity exhibited a negative correlation with fake reviews, with coefficients of -0.1740 and -0.2274, respectively. This might be because fake reviews tend to be more extreme and highly subjective. The number of photos showed a positive correlation (coefficient: 0.0801) with fake reviews, which our team speculates might be due to real users generally being less inclined to share photos in their reviews, while fake reviews aim to persuade consumers by providing numerous images to "prove" the product's quality. For example, in real-world situations, merchants often provide consumers with coupons to encourage positive reviews and offer an even higher discount when requesting a photo or video review (Fera). As for the positive correlation between the year and fake reviews (coefficient: 0.0283), this could be attributed to the increasing prevalence of fake reviews over time, driven by advances in technology that lower the cost of generating them.

Predictive Model

Since we are dealing with text data and need to perform text analysis, we initially considered three models: Artificial Neural Networks (ANN), Recurrent Neural Networks (RNN), and Support Vector Machines (SVC). ANN and RNN are well-suited for natural language processing tasks, while SVC is highly effective for common text classification problems. Consequently, our team decided to train four models—ANN, RNN, and SVC—to ensure optimal performance and results.

During the research process, we discovered that the Artificial Neural Network (ANN) model outperformed the other two models in terms of its F1 score, which demonstrated the best overall performance. As a result, we decided to focus our efforts on further studying and fine-tuning the ANN model and its parameters to optimize its performance, without the need to delve into too many specific numerical details.

Initially, we did not adjust the sample weights and ratios; we simply cleaned the data and trained the models, resulting in a low F1 score of approximately 0.12. After resampling the data, we observed a significant improvement in the F1 score, which increased to 0.22. Further improvements were achieved by implementing word embeddings and reweighting, which raised the F1 score to 0.28. Subsequently, we performed targeted feature engineering on the samples, extracting year, month, and day information from the data and identifying polarity and subjectivity as additional features. After fine-tuning the model parameters, we were able to achieve a final F1 score of 0.39. The whole model design process is illustrated clearly and concisely (Figure 3).

Results

Our in-depth analysis and predictive model development for RIT-Commerce revealed key insights and outcomes. We identified a significant data imbalance and determined critical features correlated with fake reviews, such as review rating, number of photos, year, month, is_weekday, polarity, and subjectivity. The Artificial Neural Network (ANN) model emerged as the best performer, and through optimization strategies such as resampling, word embeddings, reweighting, targeted feature engineering, and parameter fine-tuning, we improved the F1 score from 0.12 to 0.39. These results demonstrate the model's effectiveness in assisting RIT-Commerce to identify and eliminate biased fake reviews, thereby enhancing the review section's authenticity and positively impacting the platform's overall performance and profitability.

Recommendations

Our recommendation for RIT-Commerce is to completely remove fake reviews from their platform. While paying actual people to generate fake positive reviews may initially increase a product's ranking and purchases, it can greatly damage the product's quality and consumer loyalty (Clark). This loss permanently reduces the brand's competitiveness and cannot be re-established by any marketing effort. Furthermore, it can create a toxic market for sellers.

To arrive at this decision, we considered several factors that are essential for a thriving e-commerce platform. After analyzing products with the most fake reviews in the dataset, we found that as the number of fake reviews increases for a seller, the sales initially increase rapidly within a few days, but then fall back to the original level. If the percentage of fake reviews is too high, then sales might even decrease after the abrupt increase because identifying fake reviews has a significant negative impact on purchasing behaviors ("Fake Online Reviews Cost $152 Billion a Year. Here's How e-Commerce Sites Can Stop Them.").

Furthermore, fake reviews unnaturally inflate product rankings. For example, product ID 445's ranking rose sharply with increased fake reviews, a trend persisting even a month later (Figure 4 & Figure 5). If most sellers buy fake reviews, rankings will depend more on these reviews than product quality, raising product costs and hurting the affordability of loyal consumers (Pitman).

Additionally, fake reviews are deemed unhelpful, as they're subjective and lack factual information. Consumers rely on facts to assess product quality and attributes, which are scarce in fake reviews. Even if fake reviews contain facts, they're likely biased. Consumers will eventually realize these reviews aren't helpful.

By covering all the above aspects, we discovered several tangible managerial implementations that RIT-Commerce could consider for further improvements in moderating the review section. They would be summarized as follows.

Implementation

It is unrealistic to completely eradicate fake reviews from a popular platform like RIT-Commerce because paying for them is a common marketing technique for independent sellers. We suggest that RIT-Commerce allows users to report reviewer IDs that have posted suspicious reviews and assign human moderators to scrutinize the reported content. Additionally, vendors found with too many fake reviews should be banned for a certain period of time since they can be harmful to the platform's reputation and image (Petts). On the seller's end, RIT-Commerce could also allow them to remove reviews or report reviewers if they found them suspicious.

In conclusion, RIT-Commerce endeavors to eliminate fabricated reviews to ensure consistency in sales for sellers, and to promote equity among all participants of the platform, thereby safeguarding the interests of both vendors and consumers. This serves as the foundation for the sound and sustainable development of widely-used online marketplaces such as RIT-Commerce.

Bibliography

Brownlee, Jason. "Overfitting and Underfitting with Machine Learning Algorithms." *MachineLearningMastery.com*, 12 Aug. 2019, https://machinelearningmastery.com/overfitting-and-underfitting-with-machine-learning-algorithms/.

Clark, Jon. "18 Online Review Statistics Every Marketer Should Know." *Search Engine Journal*, 13 Jan. 2023, https://www.searchenginejournal.com/online-review-statistics/329701/#close.

"Fake Online Reviews Cost $152 Billion a Year. Here's How e-Commerce Sites Can Stop Them." *World Economic Forum*, https://www.weforum.org/agenda/2021/08/fake-online-reviews-are-a-152-billion-problem-heres-how-to-silence-them/.

Fera Team. "Offering Incentives for Reviews: Can I Offer a Discount or Loyalty Points for Leaving a Review?" *Fera*, https://www.fera.ai/blog/posts/offering-incentives-for-reviews-can-i-offer-a-discount-or-loyalty-points-for-leaving-a-review.

Pitman, Jamie. "Fake Reviews Are a Real Problem: 8 Statistics That Show Why." *BrightLocal*, 9 Sept. 2022, https://www.brightlocal.com/learn/review-management/fake-reviews/why-are-fake-reviews-a-problem/.

Petts, Bailey. "Should You Be Using Fake Reviews? (Here's Why They're Dangerous)." *Cloutly*, Blog | Cloutly | Reviews, Reputation and Digital Marketing Insights, 1 Nov. 2022, https://cloutly.com/blog/using-fake-reviews-dangerous/.

Appendix

Figure 1

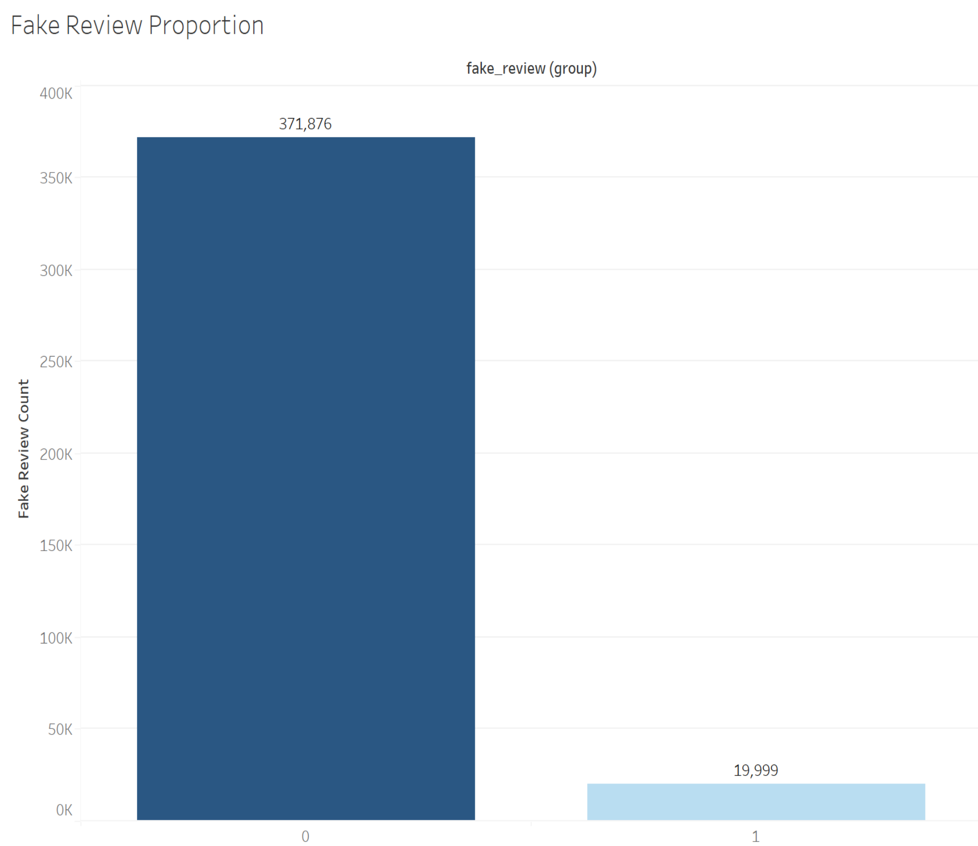

Figure 2

Model:	Logit	Pseudo R-squared:	0.095
Dependent Variable:	fake_review	AIC:	143028.4672
Date:	2023-03-25 16:11	BIC:	143148.1329
No. Observations:	391875	Log-Likelihood:	-71503.
Df Model:	10	LL-Null:	-78982.
Df Residuals:	391864	LLR p-value:	0.0000
Converged:	1.0000	Scale:	1.0000
No. Iterations:	9.0000		

	Coef.	**Std.Err.**	**z**	**P>|z|**	**[0.025**	**0.975]**
const	-60.3065	9.3899	-6.4225	0.0000	-78.7102	-41.9027
review_rating	0.0367	0.0066	5.5595	0.0000	0.0238	0.0497
number_of_photos	0.0801	0.0137	5.8643	0.0000	0.0533	0.1069
helpful_vote	0.0012	0.0003	4.3769	0.0000	0.0007	0.0017
year	0.0283	0.0047	6.0898	0.0000	0.0192	0.0374
month	-0.0467	0.0023	-20.2493	0.0000	-0.0513	-0.0422
day	-0.0006	0.0009	-0.7171	0.4733	-0.0023	0.0011
is_weekday	0.0311	0.0171	1.8251	0.0680	-0.0023	0.0646
polarity	-0.1740	0.0344	-5.0645	0.0000	-0.2414	-0.1067
subjectivity	-0.2274	0.0405	-5.6161	0.0000	-0.3068	-0.1481
labeled_product_id	0.0047	0.0000	108.4792	0.0000	0.0046	0.0048

Figure 3

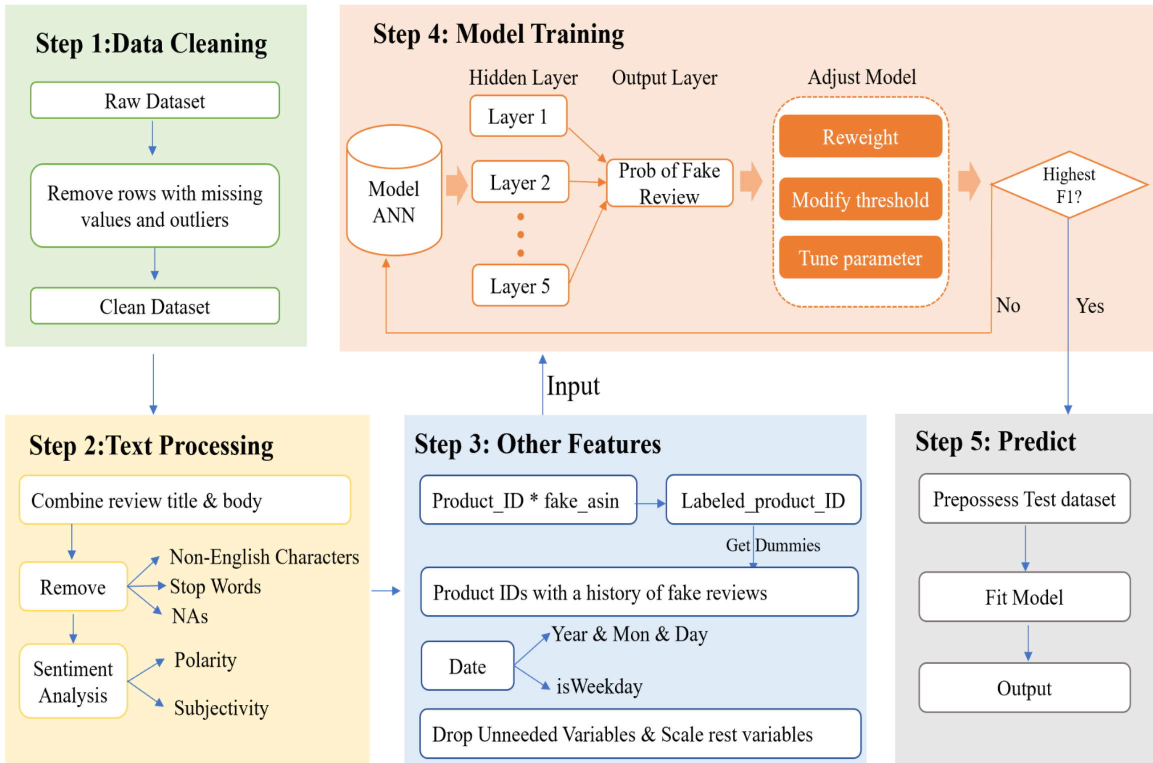

Step 1:Data Cleaning
Raw Dataset
Remove rows with missing values and outliers
Clean Dataset
Step 4: Model Training
Hidden Layer
Output Layer
Adjust Model
Model ANN
Layer 1
Layer 2
Layer 5
Prob of Fake Review
Reweight
Modify threshold
Tune parameter
Highest F1?
No
Yes
Input
Step 2:Text Processing
Combine review title & body
Remove
Non-English Characters
Stop Words
NAs
Sentiment Analysis
Polarity
Subjectivity
Step 3: Other Features
Product_ID * fake_asin
Labeled_product_ID
Get Dummies
Product IDs with a history of fake reviews
Date
Year & Mon & Day
isWeekday
Drop Unneeded Variables & Scale rest variables
Step 5: Predict
Prepossess Test dataset
Fit Model
Output

Figure 4

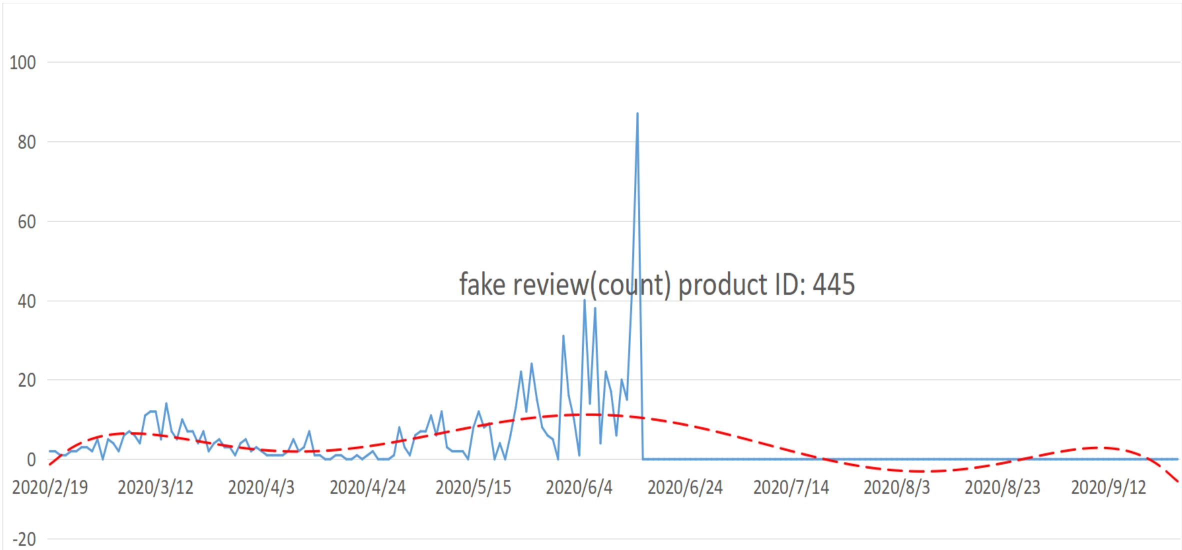

Figure 5